



US009449050B1

(12) **United States Patent**
Molina et al.

(10) **Patent No.:** **US 9,449,050 B1**
(45) **Date of Patent:** ***Sep. 20, 2016**

(54) **IDENTIFYING RELEVANT MESSAGES IN A CONVERSATION GRAPH**

(71) Applicant: **Twitter, Inc.**, San Francisco, CA (US)

(72) Inventors: **Marcel Molina**, San Francisco, CA (US); **Ross Cohen**, San Francisco, CA (US); **Kyle Maxwell**, San Francisco, CA (US); **Stuart Hood**, San Francisco, CA (US); **Cara Meverden**, San Francisco, CA (US); **Coleen Baik**, San Francisco, CA (US)

(73) Assignee: **Twitter, Inc.**, San Francisco, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 426 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **13/770,955**

(22) Filed: **Feb. 19, 2013**

Related U.S. Application Data

(63) Continuation-in-part of application No. 13/736,634, filed on Jan. 8, 2013, now Pat. No. 9,143,468.

(51) **Int. Cl.**

G06F 15/16 (2006.01)

G06F 17/30 (2006.01)

(52) **U.S. Cl.**

CPC **G06F 17/3053** (2013.01)

(58) **Field of Classification Search**

USPC 709/206, 204; 705/756, 764, 709, 751, 705/745

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,275,811 B1 * 8/2001 Ginn G06Q 30/02 705/14.31

7,546,131 B1 6/2009 Sidi et al.

7,752,649 B1 7/2010 Harvey et al.
7,830,925 B1 11/2010 Harvey et al.
2006/0200442 A1 * 9/2006 Parikh G06F 17/30648
2008/0012701 A1 * 1/2008 Kass A61B 5/0002 340/539.11
2009/0313346 A1 * 12/2009 Sood G06Q 10/10 709/207
2010/0273443 A1 * 10/2010 Forutanpour et al. 455/404.1
2012/0271860 A1 * 10/2012 Graham, Jr. G06Q 30/02 707/798
2013/0036177 A1 * 2/2013 Leeder et al. 709/206
2013/0110565 A1 * 5/2013 Means, Jr. G06Q 10/063 705/7.11
2013/0304726 A1 * 11/2013 Sandulescu et al. 707/722
2014/0129936 A1 * 5/2014 Richards et al. 715/716

OTHER PUBLICATIONS

Notice of Allowance for U.S. Appl. No. 14/020,447, mailed on Apr. 20, 2016, 14 pages.

* cited by examiner

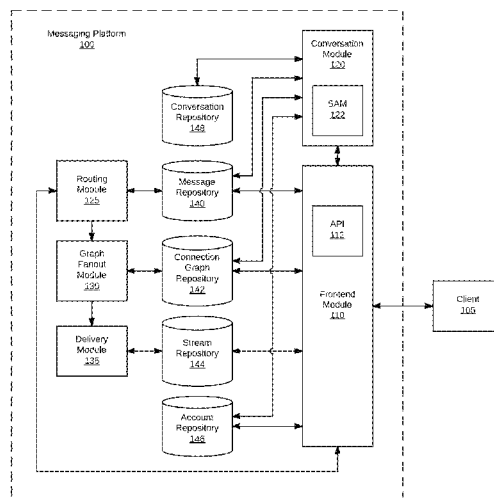
Primary Examiner — Jude Jean Gilles

(74) *Attorney, Agent, or Firm* — Brakes Hughes Bellermann LLP

(57) **ABSTRACT**

A method, system, and non-transitory computer readable medium for identifying a relevant conversation. The method can include: identifying a conversation graph including identifiers of messages posted by a set of authoring accounts; identifying a relevance score for each authoring account; applying, by a computer processor, a first relevance filter to the conversation graph in order to modify the set of relevance scores; applying, by the computer processor and after applying the first relevance filter, a second relevance filter to the conversation graph in order to further modify the set of relevance scores; selecting, based on the set of relevance scores and after applying the second relevance filter, a subset of messages of the conversation graph for inclusion in the relevant conversation; and providing at least a portion of the relevant conversation to a client device.

26 Claims, 10 Drawing Sheets



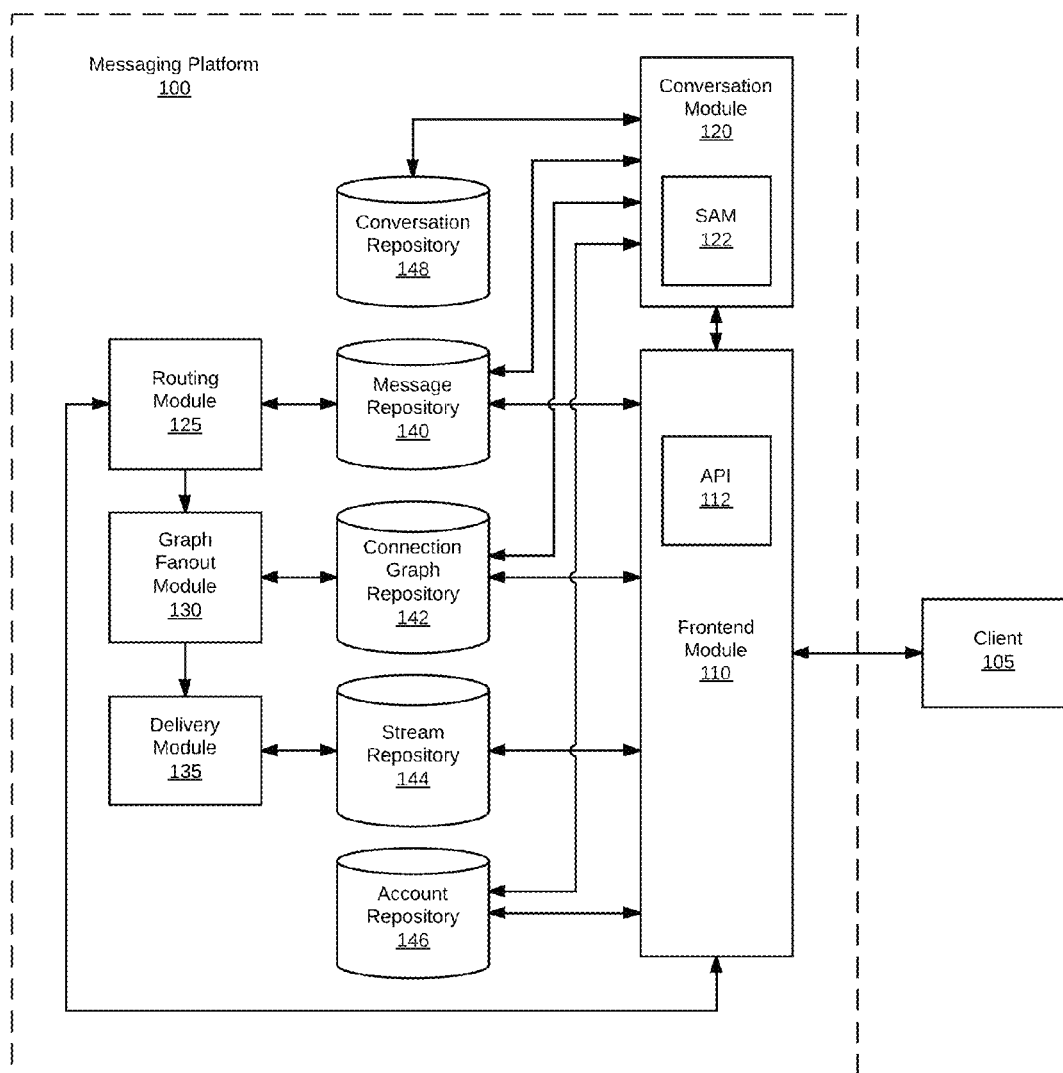


Figure 1

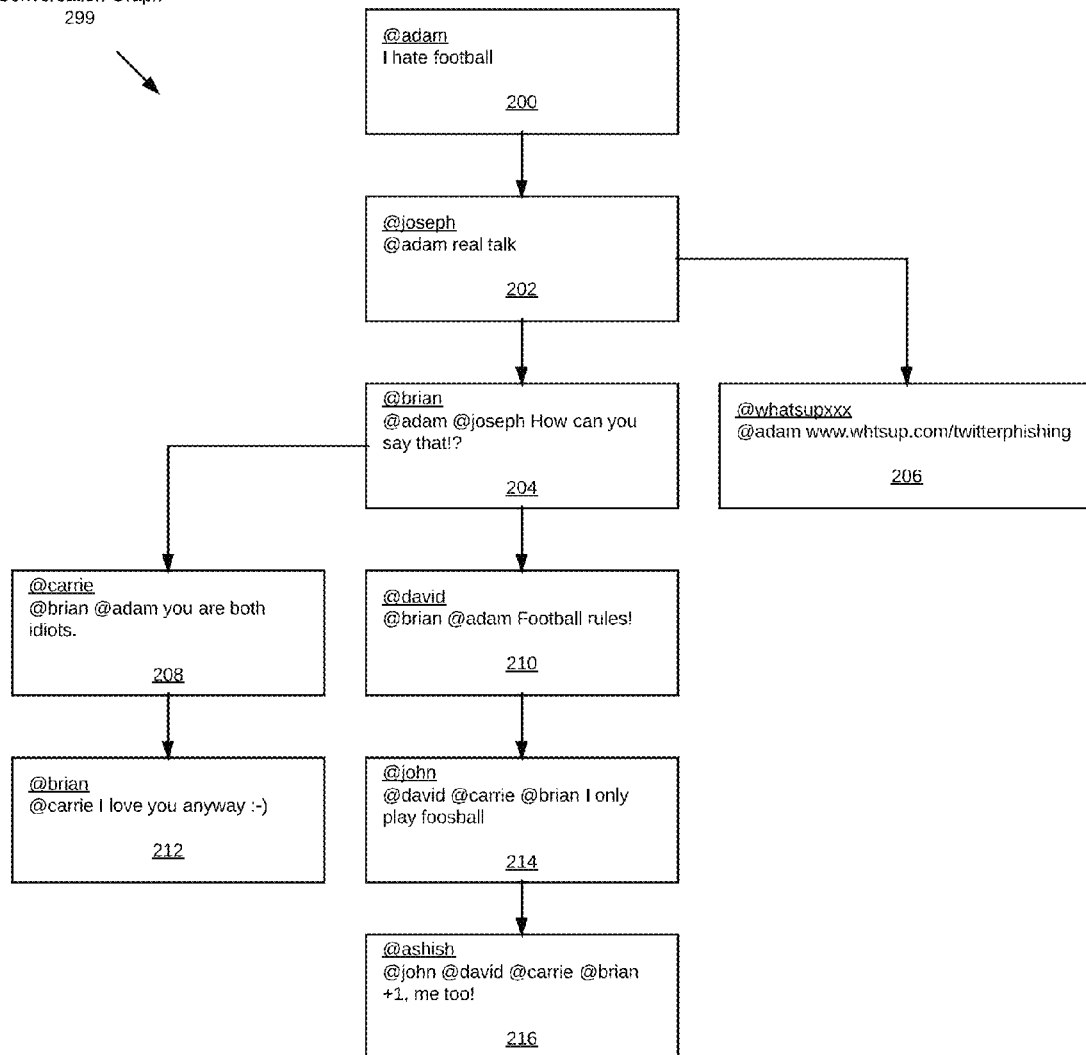

Conversation Graph
299

Figure 2

Message Table
300



*	Message ID 302	Author Account ID 304	Message Content 306	Relevant Participant(s) Account ID 308	In Reply to 310
312	1	@adam	I hate football	none	none
314	2	@joseph	@adam real talk	@adam	1
316	3	@brian	@adam @joseph How can you say that?	@adam, @joseph	2
...

Figure 3

Account Data Table
400

*	Account ID 402	Display Name 404	Followed By 406
410	@keith	Keith Jones	@brian
412	@adam	Adam Doe	@brian, @carrie, @david, @joseph
414	@carrie	Carrie Barnes	@brian, @adam
416	@david	David Swenson	@brian, @john
418	@brian	Brian James	@david, @adam, @carrie, @keith
420	@whatsupxxx	What's Up Marketing, LLC	@adam, @george, @paul
422	@joseph	Joseph Lauer	@adam, @ashish, @brian
424	@ashish	Ashish Pahal	@john, @carrie

Figure 4

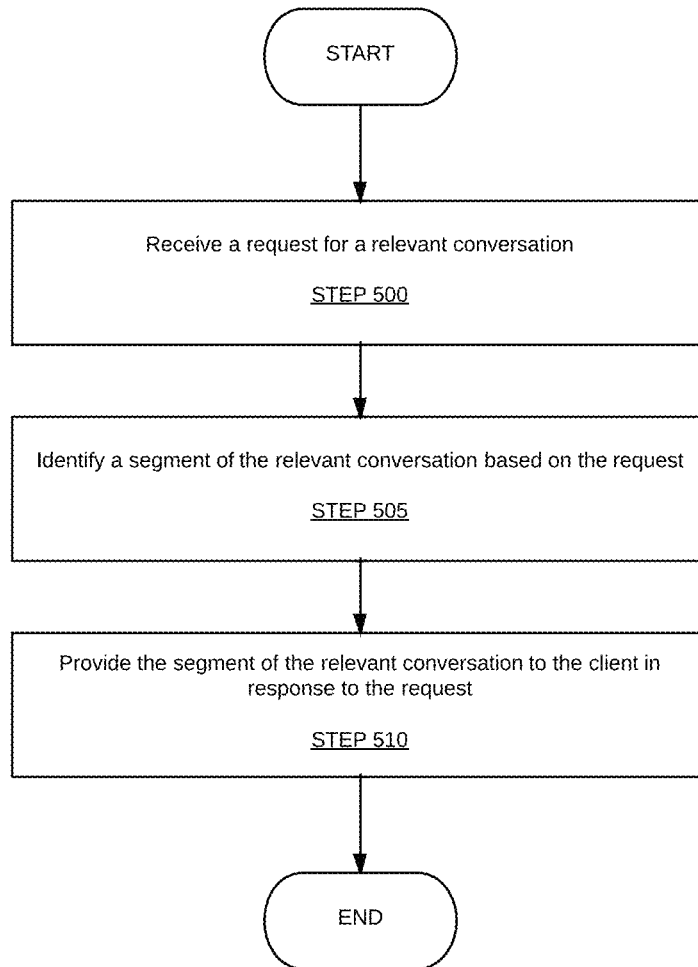


Figure 5

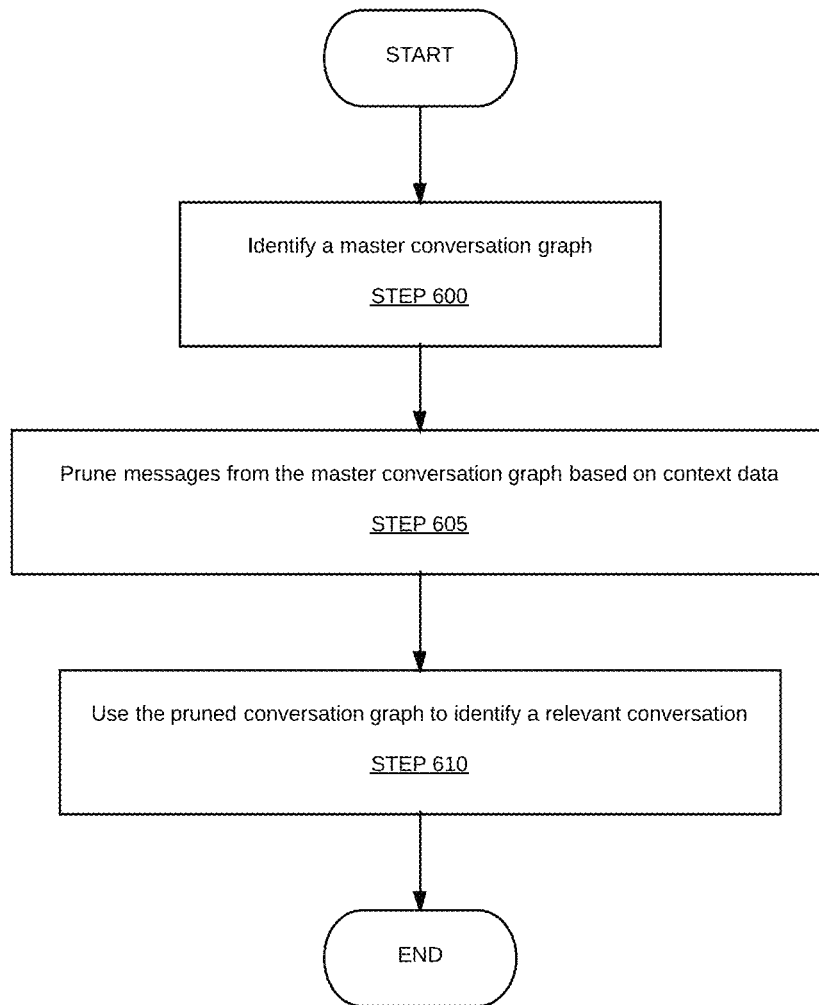


Figure 6

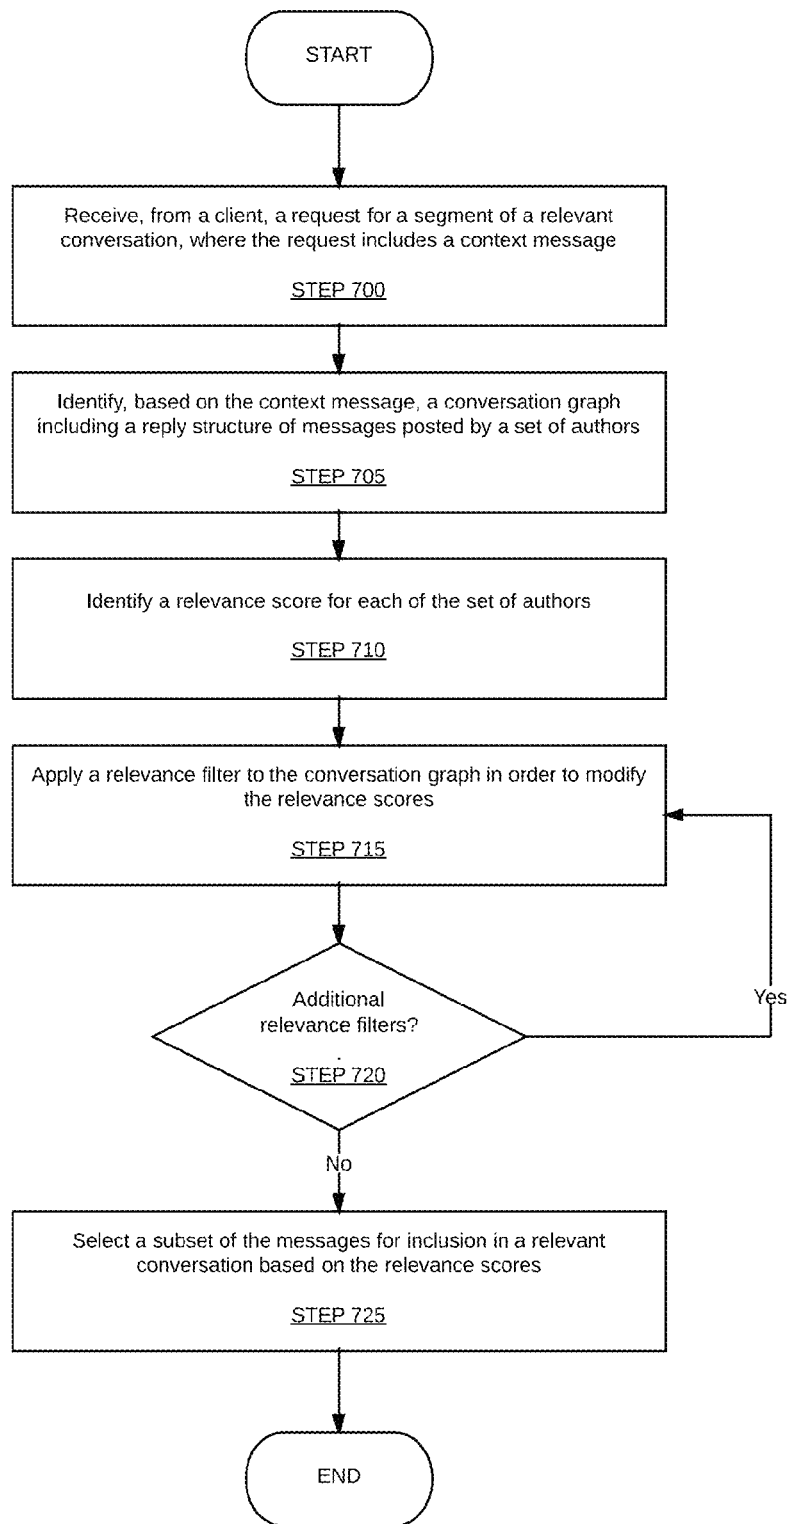


Figure 7

Mobile Device
899

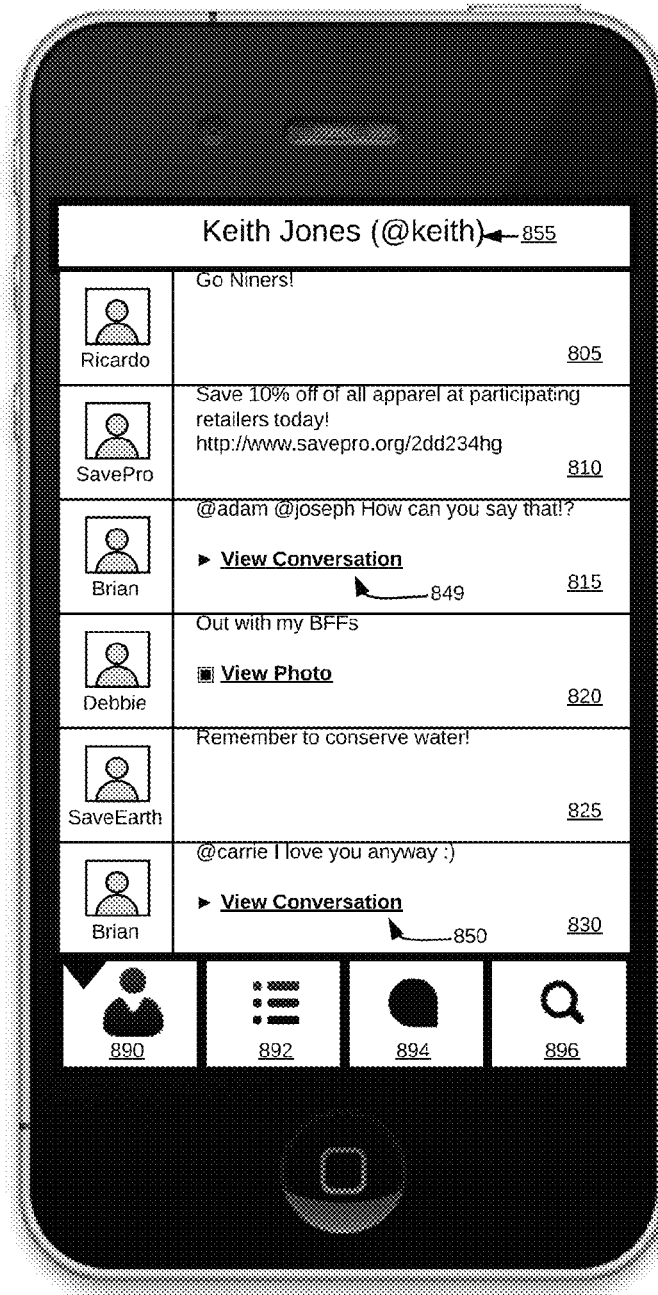


Figure 8A

Mobile Device
899

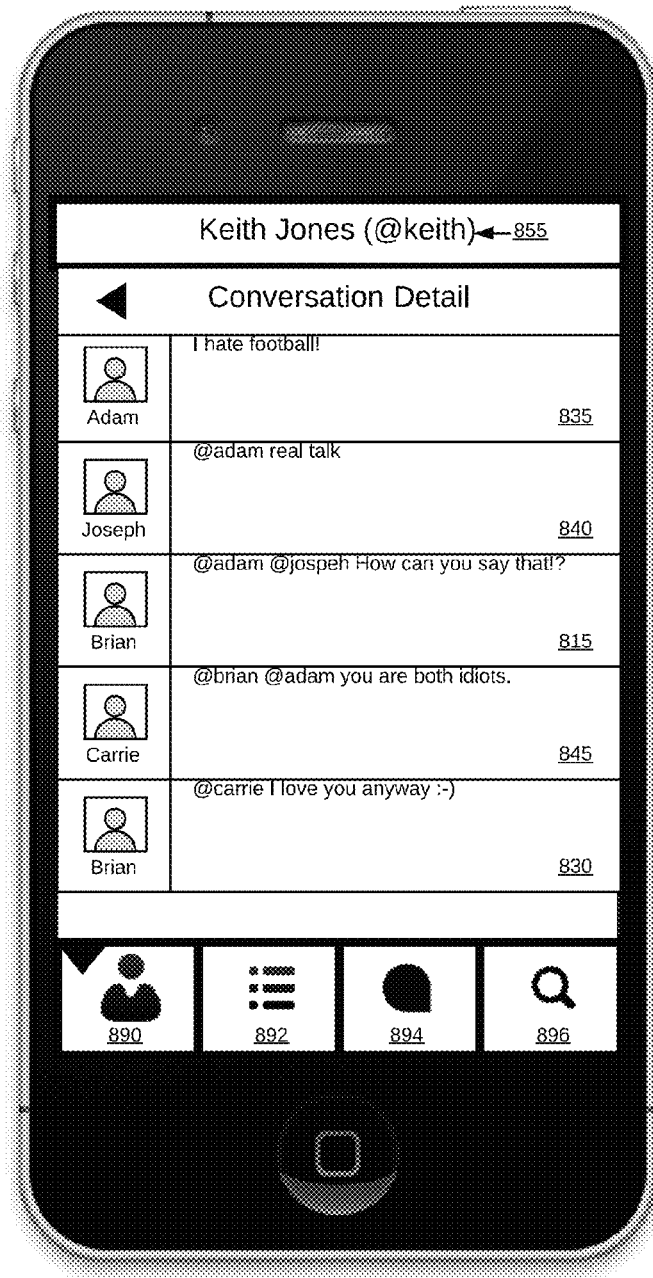


Figure 8B

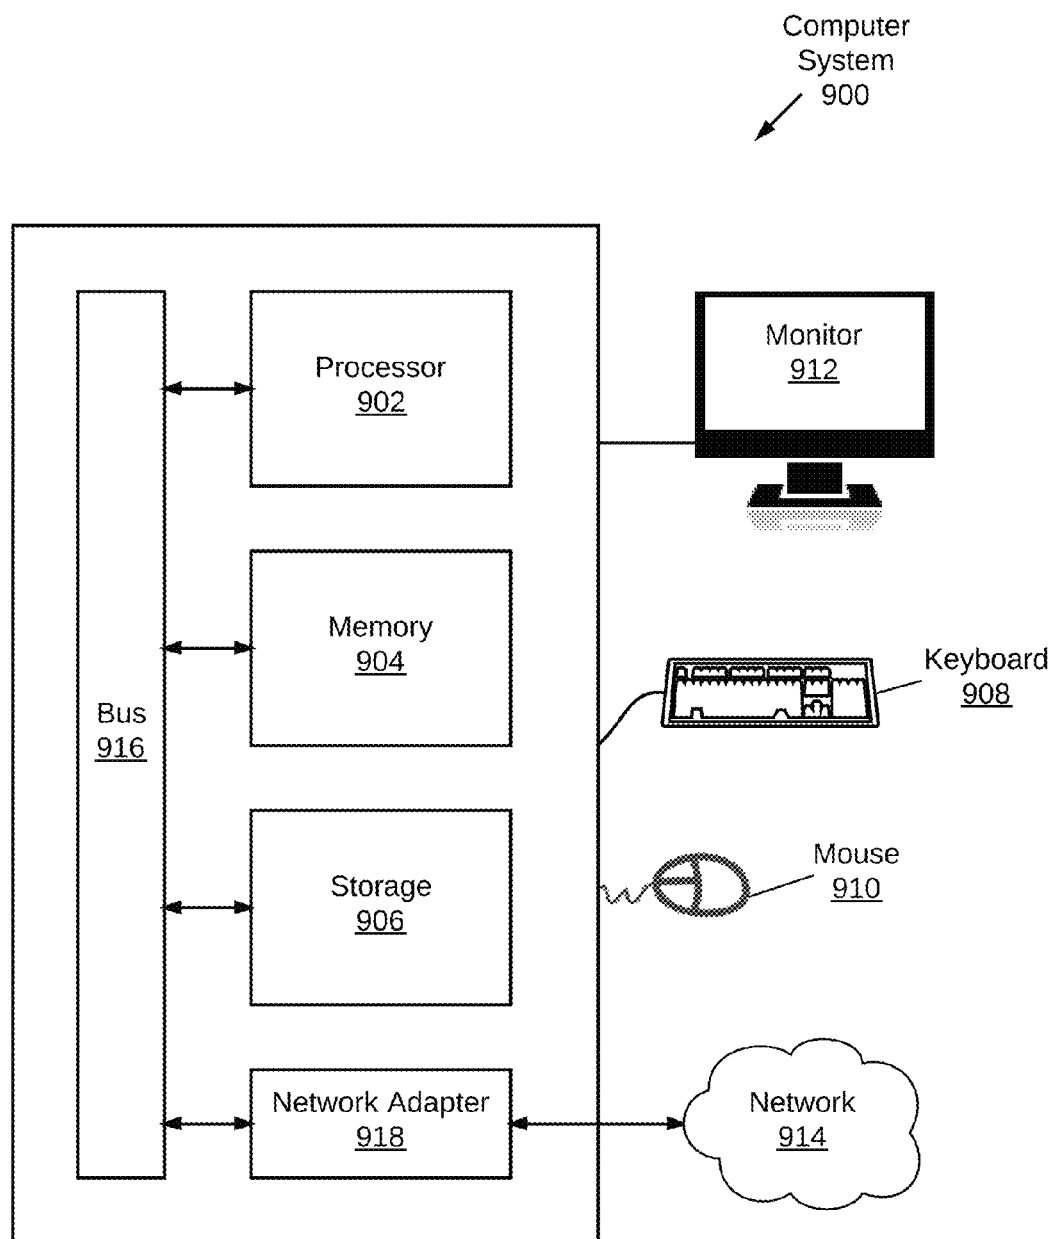


Figure 9

1

IDENTIFYING RELEVANT MESSAGES IN A CONVERSATION GRAPH

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a continuation-in-part of U.S. patent application Ser. No. 13/736,634, filed on Jan. 8, 2013 and entitled "Identifying Relevant Messages in a Conversation Graph." Accordingly, this application claims benefit from U.S. patent application Ser. No. 13/736,634 under 35 U.S.C. §120. U.S. patent application Ser. No. 13/736,634 is incorporated by reference herein, in its entirety, for all purposes.

Where a definition or use of a term in a reference which is incorporated by reference herein is inconsistent or contrary to the definition of that term provided herein, the definition of that term provided herein applies and the definition of that term in the incorporated reference does not apply.

BACKGROUND

Social networking services allow users to author and share messages with multiple users. For example, a first user can author and publish a message for other users to read. The messages can be any of a variety of lengths which may be limited by a specific messaging system or protocol. For example, a short messaging service protocol typically limits messages to 140 characters in length.

Users interested in viewing messages authored by a particular user can choose to follow the particular user. A first user can follow a second user by identifying the second user as a user the first user would like to follow. After the first user has indicated that they would like to follow the second user, the first user will be provided with messages authored by the second user. Users can choose to follow multiple users. Users can also respond to messages and thereby have conversations with one another. Both unidirectional (e.g., follower/followee) and bidirectional (e.g., friendship) relationships may exist, depending on the type of network.

A conversation can include a root message and any number of messages directly or indirectly in reply to the root message. A conversation graph is a data structure which represents two or more messages in a conversation as nodes, with edges connecting the nodes based on a reply structure of the conversation. Users can view a conversation based on messages included in the conversation graph. However, when a conversation includes a large number of messages, the number of messages displayed can be very large and may overwhelm the user.

SUMMARY

In general, in one aspect, the invention relates to a method for identifying a relevant conversation. The method can include: identifying a conversation graph including identifiers of messages posted by a set of authoring accounts; identifying a set of relevance scores where each of the set of relevance scores corresponds to an authoring account of the set of authoring accounts; applying, by a computer processor, a first relevance filter to the conversation graph in order to modify the set of relevance scores; applying, by the computer processor and after applying the first relevance filter, a second relevance filter to the conversation graph in order to further modify the set of relevance scores; selecting, based on the set of relevance scores and after applying the

2

second relevance filter, a subset of messages of the conversation graph for inclusion in the relevant conversation; and providing at least a portion of the relevant conversation to a client device.

In general, in one aspect, the invention relates to a system for identifying a relevant conversation. The system can include: a computer processor; a conversation module executing on the computer processor and configured to: identify a conversation graph including identifiers of messages posted by a set of authoring accounts, identify a set of relevance scores where each of the set of relevance scores corresponds to an authoring account of the set of authoring accounts, apply a first relevance filter to the conversation graph in order to modify the set of relevance scores, apply, after applying the first relevance filter, a second relevance filter to the conversation graph in order to further modify the set of relevance scores, select, based on the set of relevance scores and after applying the second relevance filter, a subset of messages of the conversation graph for inclusion in the relevant conversation, and provide at least a portion of the relevant conversation to a client device.

In general, in one aspect, the invention relates to a non-transitory computer readable medium including instructions for identifying a relevant conversation. The instructions, when executed by a computer processor, can include functionality to: identify a conversation graph including identifiers of messages posted by a set of authoring accounts; identify a set of relevance scores where each of the set of relevance scores corresponds to an authoring account of the set of authoring accounts; apply a first relevance filter to the conversation graph in order to modify the set of relevance scores; apply, after applying the first relevance filter, a second relevance filter to the conversation graph in order to further modify the set of relevance scores; select, based on the set of relevance scores and after applying the second relevance filter, a subset of messages of the conversation graph for inclusion in the relevant conversation; and provide at least a portion of the relevant conversation to a client device.

Other aspects of the invention will be apparent from the following description and the appended claims.

BRIEF DESCRIPTION OF DRAWINGS

FIG. 1 depicts a diagram of a messaging platform in accordance with one or more embodiments of the invention.

FIG. 2 depicts a graphical representation of a conversation graph, in accordance with one or more embodiments of the invention.

FIG. 3 depicts a message table, in accordance with one or more embodiments of the invention.

FIG. 4 depicts an account data table, in accordance with one or more embodiments of the invention.

FIGS. 5-7 depict flowcharts of methods for identifying relevant messages, in accordance with one or more embodiments of the invention.

FIGS. 8A and 8B depict example user interfaces displayed by a mobile device, in accordance with one or more embodiments of the invention.

FIG. 9 depicts a high-level block diagram of a computer for identifying relevant messages in a conversation graph, in accordance with one or more embodiments of the invention.

DETAILED DESCRIPTION

Specific embodiments of the invention will now be described in detail with reference to the accompanying

figures. Like elements in the various figures are denoted by like reference numerals for consistency.

In the following detailed description of embodiments of the invention, numerous specific details are set forth in order to provide a more thorough understanding of the invention. However, it will be apparent to one of ordinary skill in the art that the invention may be practiced without these specific details. In other instances, well-known features have not been described in detail to avoid unnecessarily complicating the description.

In general, embodiments of the invention provide a method and system for identifying relevant messages in a conversation graph. A conversation graph is identified including messages posted by a set of authors. A set of relevance scores associated with the messages are identified. Relevance filters are applied to the conversation graph in order to modify the set of relevance scores. A subset of the messages are then selected for inclusion in a relevant conversation based on the set of relevance scores. At least a portion of the relevant conversation is provided to a client in response to a request.

FIG. 1 shows a messaging platform (100) and a client (105), in accordance with one or more embodiments of the invention. As shown in FIG. 1, the messaging platform (100) can have multiple components including a frontend module (110), an application programming interface (API 112), a conversation module (120), a semantic analysis module (SAM 122), a conversation repository (148), a routing module (125), a graph fanout module (130), a delivery module (135), a message repository (140), a connection graph repository (142), a stream repository (144), and an account repository (146). Various components of the messaging platform (100) of FIG. 1 may be located on the same device (e.g., a server, mainframe, desktop Personal Computer (PC), laptop, mobile phone, kiosk, cable box, and any other device) or may be located on separate devices connected by a network (e.g., a local area network (LAN), the Internet, etc.). Those skilled in the art will appreciate that there may be more than one of each separate component running on a device, as well as any combination of these components within a given embodiment of the invention.

In one or more embodiments of the invention, the messaging platform (100) is a platform for facilitating real-time communication between one or more entities. For example, the messaging platform (100) may store millions of accounts of individuals, businesses, and/or other entities (e.g., pseudonym accounts, novelty accounts, etc.). One or more users of each account may use the messaging platform (100) to send messages to other accounts inside and/or outside of the messaging platform (100). The messaging platform (100) may be configured to enable users to communicate in “real-time”, i.e., to converse with other users with a minimal delay and to conduct a conversation with one or more other users during simultaneous sessions. In other words, the messaging platform (100) may allow a user to broadcast messages and may display the messages to one or more other users within a reasonable time frame so as to facilitate a live conversation between the users. Recipients of a message may have a predefined graph relationship with an account of the user broadcasting the message. In one or more embodiments of the invention, the user is not an account holder or is not logged in to an account of the messaging platform (100). In this case, the messaging platform (100) may be configured to allow the user to broadcast messages and/or to utilize other functionality of the messaging platform (100) by associating the user with a temporary account or identifier.

For purposes of this disclosure, a conversation graph is a data structure or a collection of data structures representing a reply structure of messages posted to the messaging platform (100). In one or more embodiments of the invention, a conversation graph can represent any subset of messages of a given conversation. In one or more embodiments of the invention, a master conversation graph can refer to a type of conversation graph which represents all messages in a given conversation (i.e., all messages within the same reply structure).

FIG. 2 is an example depiction of a conversation graph (299) comprising messages 200 through 216. A conversation graph is a data structure including messages of a conversation. In one or more embodiments, the conversation graph includes a root message (e.g., message 200 of FIG. 2) and one or more messages connected, either directly or indirectly, to the root message (e.g., messages 204-216 of FIG. 2). The relationships between the nodes of the conversation graph represent a reply structure of the conversation, wherein each child node is a reply message to at least one corresponding parent node, in accordance with various embodiments of the invention. The author of a message is shown underlined in FIG. 2. For example, the author of message 200 is associated with the account “@adam” as identified by the underlining. Messages 200-216 correspond to message data stored in table 300 of FIG. 3. For example, record 312 of table 300 is identified in message ID field 302 as message “1” which corresponds to message 200 depicted in FIG. 2. Record 314 of table 300 is identified in message ID field 302 as message “2” which corresponds to message 202 in FIG. 2. Record 316 of table 300 is identified in message ID field 320 as message “3” which corresponds to message 204 in FIG. 2.

Although conversation graph 299 is depicted in FIG. 2 using a graphical format, the messages depicted in FIG. 2, in one or more embodiments, can be stored in a message repository (e.g., message repository 140 of FIG. 1) in any appropriate data structure. In this example, message 200 is a root message (i.e., the parent of all messages shown in FIG. 2). Child message 202 is a reply message that is in reply to message 200 as identified by the line connecting message 202 to message 200. Child messages 204 and 206 are reply messages that are in reply to message 202, the relationships of which are shown by the lines connecting the messages. In this example, messages 202-216 are reply messages under root message 200.

Conversation graph 299, in one or more embodiments, is constructed based on the reply structure of messages 200-216. The reply structure may be identified based on metadata associated with each message and/or reply information identified from within the message content, in accordance with various embodiments. For example, in one embodiment, metadata stored in the In Reply to field 210 of table 200 is used to generate a conversation graph. More specifically, the metadata maintains the reply relationships between multiple messages of a conversation. Thus, in this example, each message can be in reply to only one parent message, and only messages which are in the reply structure of the root message are included in the conversation graph. It should be noted that displaying messages solely from accounts having a predefined graph relationship with a user’s account may render the conversation incomprehensible to the user.

It should be noted that messages can be categorized into various types. For example, a message posted by a particular author without an indication that the message is directed to a particular account may be referred to, in one embodiment,

5

as a broadcast message or simply a message. Broadcast messages are displayed to users with accounts associated with (e.g., following, friends with, etc.) the account of the broadcast message. A message in reply to another message is considered a reply message or, more simply, a reply. In one embodiment, the category of a message determines whether the message will be designated to be displayed to another account or included in other account's message streams. A direct message (also referred to as a private message) is a message that explicitly identifies one or more accounts as recipients. For example, a message from a user associated with the account @keith that is explicitly directed (by the user authoring the message) to the account @carrie is considered a direct message from @keith to @carrie.

FIG. 3 depicts an example of a message table 300 which includes multiple records (e.g., 312, 314, 316) containing message data fields 302-310 associated with a message identified by message ID 302. Author Account ID field 304 contains an identifier indicating an account associated with an author of the message identified by message ID 302. In one or more embodiments, each account is associated with a unique account ID. For example, a user named Adam Smith may select the identifier "adam", which is represented as "@adam". Other identifiers, such as alpha-numeric identifiers, can be selected as well. Message content 306 contains the content of a message identified by message ID 302 field. Typically, the content of a message consists of text. However, in one or more embodiments of the invention, message content can contain various media and/or may include one or more references to various media (e.g., a uniform resource locator (URL) of a video).

In one or more embodiments, relevant participant(s) account ID 308 field contains identifiers of one or more relevant accounts. In one embodiment, the relevant participant(s) account ID 308 excludes the authoring account and/or the parent account (i.e., the author of a message which the message is in reply to) if the message is a reply message. For example, message content 306 can contain text such as "@johnqpublic is correct!". The character string "@johnqpublic" indicates that an account associated with the "@johnqpublic" identifier is mentioned in the content of the message. The mentioned account identifier ("@johnqpublic") is inserted into relevant participant(s) 308 field. Identified relevant participants can then be used as a factor in determining which messages in a conversation graph to surface to one or more users (e.g., which messages to include in a relevant conversation).

In one or more embodiments of the invention, in reply to field 310 identifies a parent message that a particular message is in reply to. For example, record 314 identifies message ID number 2 which is identified as being in reply to message ID number 1. The relationships among messages may be used to identify various messages for inclusion in a conversation graph. The reply relationships between two or more messages may be identified based on a variety of different data, in accordance with various embodiments. In one or more embodiments, the reply structure is identified based on metadata associated with each message which is received from a client used to compose the message. For example, a user may click on a "Reply" link displayed below a message displayed in the user's home stream. The client may then display a message composition box for drafting a reply message. The client may submit metadata including the reply relationship (i.e., a message ID of the parent message) with the reply message. Conversely, in another example, the reply relationship may be explicitly defined by the user within the message content. In this example, the

6

reply structure may be identified by identifying one or more account IDs and/or message IDs mentioned within the message content. The designation between authors of parent messages and relevant participants, if any, may be made based on any number of syntactical rules. Thus, the message platform 100 may require that if a message begins with an account ID (e.g., "@john"), the message will be marked as being in reply to a latest message from the mentioned account ID. Thus, the reply to field 310 may be populated based on any means of identifying a reply structure of two or more messages. Lastly, it should be noted that while many of the examples depicted herein include a requirement that each message can only have one parent message, embodiments of the invention may be implemented in a system allowing multi-parent reply structures (i.e., a single message being in reply to two or more parent messages).

FIG. 4 depicts account data table 400 which includes multiple records 410-424 containing user data 404-406 associated with an account identified by account ID 402. Account ID field 402 contains an identifier which uniquely identifies each account of the messaging platform (100). In one embodiment, the account ID is selected by a user. Although many users may choose to select an account ID that is a concatenation of the user's first and last name, or the name of a business or service, the account ID can be any character string (e.g., an alpha-numeric string). Display name 404 contains an identifier indicating a user's selected display name. In one or more embodiments, display name 404 is a character string, typically chosen by a user as a name used by the messaging platform (100) for purposes of display. For example, display name 404 may be a business name, a pseudonym, and/or any other display name selected by a user of the account. Each account ID 402 can be associated with a display name 404, although display names are not required to be unique. In one or more embodiments, user names and account IDs are used to facilitate recognition of authors by readers and are used by messaging platform 100 to perform various functions.

In one or more embodiments, followed by field 406 contains identifiers indicating accounts following an account identified by account ID 402. In one or more embodiments, a user can choose to follow another account. A user following a particular account receives messages authored by a user associated with the particular account intended for public viewing. For example, if the user associated with the particular account enters a message for public viewing, all users following the particular account can view the message. In one or more embodiments, users following the particular account will receive a push notification (on their computing device) indicating that a user associated with the particular account has posted a message. In one or more embodiments, users following the particular account will receive a copy of the message the user of the particular account authored for public viewing. For example, FIG. 4 shows that account "@adam" is followed by "@brian", "@carrie", "@david", and "@joseph". As such, "@brian", "@carrie", "@david", and "@joseph" are subscribed to receive messages authored by a user associated with the account "@adam". In one or more embodiments, followed by field 406 is implemented as a bidirectional relationship field (e.g., friendship) and represents a mutual connection between accounts of the messaging platform.

Returning to FIG. 1, in one or more embodiments of the invention, the conversation module (120) includes functionality to identify one or more relevant conversations based on a conversation graph. For example, the conversation module (120) may be configured to perform the functions described

herein in connection with FIGS. 5-7 and/or to receive requests from a user via the front end module (110). In one or more embodiments of the invention, relevant conversations determined by the conversation module (120) are stored in the conversation repository (148). The conversation repository (148) includes, in one or more embodiments, functionality to transmit one or more relevant conversations, and/or messages from a relevant conversation, in response to requests from, for example, the conversation module (120) or the front end module (110) via the conversation module (120).

In one or more embodiments of the invention, the conversation module (120) includes functionality to identify a conversation graph. The conversation graph can be retrieved, for example, from the conversation repository (148). Conversation graphs can be constructed based on messages in the message repository (140) and stored in the conversation repository (148) by the conversation module (120) or one or more related components of the messaging platform (100), in accordance with various embodiments.

In one or more embodiments of the invention, the conversation module (120) includes functionality to receive a request for a relevant conversation. The request can be for a portion of the relevant conversation, such as, for example, the 10 nearest relevant messages to a context message. Any other criteria or data can be provided, in conjunction with the request, to identify at least a segment of the relevant conversation. For example, the request can include a parameter value defining a number of requested messages. The conversation module (120) can be configured to return the number of requested messages from a relevant conversation based on the parameter value.

In one or more embodiments of the invention, the request for a relevant conversation includes context data. One example of context data is a context message. A context message can be any message of the conversation which is identified by the request. The context message can be used to identify a context of a request for a relevant messages. Any one of the messages in a conversation graph can be identified as a context message. Examples of a context message can include, but are not limited to, a message selected by a user, a message including a specified keyword or topic, a message returned in response to a search request, a message associated with an advertisement, and/or any other message included in at least one conversation graph. In one or more embodiments, the context message is identified based on a user selecting a message from a message stream presented via front end module 110. In other embodiments, the context message may be identified as a message including a specified keyword or topic, a message returned in response to a search request, and/or a message associated with an advertisement. Other examples of context data can include, but are not limited to, a number of requested messages, a pagination offset indicating a window of messages (e.g., a position below a context message), a requesting account (i.e., a requester of the relevant conversation), one or more relevant topics (i.e., predefined categories used to adjust message/account relevance scores, discussed below), and any other information relevant to identifying messages within the conversation graph.

In one or more embodiments of the invention, the conversation module (120) includes functionality to prune a conversation graph based on context data. In one or more embodiments, the conversation graph can be pruned prior to applying any relevance filters, to include only (1) the root message, (2) the context message, (3) all messages connecting the root message to the context message, and/or (4) all

direct or indirect children of the context message. In this way, for example, the conversation graph can be modified to include only a single branch of messages connecting the root message to the context message, as well as any number of branches below the context message.

In one or more embodiments of the invention, the conversation module (120) includes functionality to assign relevance scores to authoring accounts of messages of a conversation graph. Thus, each authoring account of a message of the conversation graph is assigned a relevance score (i.e., an account relevance score). The relevance score can be an arbitrary initial score or any other baseline score assigned prior to applying one or more relevance filters. For example, all account relevance scores can be initialized to a value of zero.

In one or more embodiments of the invention, the conversation module (120) includes functionality to apply relevance filters to the conversation graph in order to modify one or more the relevance scores of the authoring accounts. Relevance filters which are used to modify relevance scores of authoring accounts can be referred to as account relevance filters.

In the following description of relevance filters, an increase in a relevance score can be implemented as a relative increase or a fixed amount increase. For example, if a relevance filter is said to increase the relevance score of an authoring account, the account relevance score can be incremented by 2, or all other relevance scores can be decremented by 2 resulting in a relative increase for the authoring account. In another example, the relevance score of a message can be increased by 20%, or based on a factor of the standard deviation of all the relevance scores. Likewise, in one or more embodiments of the invention, a decrease in a relevance score can be implemented as a relative decrease or a fixed amount decrease.

In one or more embodiments of the invention, the conversation module (120) includes functionality to apply a relevance filter to increase the relevance score(s) of an authoring account of a context message and/or an authoring account of a root message of the conversation graph. In one or more embodiments, even if the authoring account(s) of the root and/or context message are not selected for inclusion in the relevant conversation, the root message and/or the context message are always included in the relevant conversation.

In one or more embodiments of the invention, applying the relevance filter includes designating one or more authoring accounts as blessed. For example, an authoring account of the context message and/or root message can be designated as blessed. The conversation module (120) can be configured to increase the relevance scores of the blessed accounts and/or to identify one or more messages of the conversation graph which are replied to by a blessed account. As part of applying the relevance filter, the conversation module (120) can then increase the relevance scores of the authoring accounts of the replied-to messages. In one or more embodiments of the invention, the relevance scores of the authoring accounts of the replied-to messages can be increased by a lower amount relative to the relevance scores of the blessed accounts. In one or more embodiments of the invention, the conversation module (120) is configured to designate authoring accounts having a relevance score exceeding a predefined threshold as blessed. By identifying blessed accounts, and then increasing the relevance scores of one or more accounts associated with the blessed accounts (e.g., based on their reply relationship in the conversation graph), the conversation module (120) can be

configured to repeat the process by checking to see if any of the modified relevance scores now exceed the predefined threshold. If so, the conversation module (120) can designate the authoring accounts with the exceeding relevance scores as blessed. The conversation module (120) can be configured to repeat this process until convergence (i.e., until no additional authoring accounts are designated as blessed). In one or more embodiments of the invention, the authoring account(s) of the context message and/or root message are not assigned relevance scores, but are always included in the relevant conversation. Whether or not the authoring account(s) of the context message and/or root message are assigned relevance scores, in one or more embodiments, the authoring account(s) of the context message and/or root message can be designated as blessed and used by the conversation module (120) for the purpose of identifying additional blessed accounts (as described).

In one or more embodiments of the invention, the conversation module (120) includes functionality to apply a relevance filter to modify the relevance scores of one or more authoring accounts based on a credibility rating of each account. As part of applying the relevance filter, the conversation module (120) can be configured to identify a credibility rating of each authoring account of the conversation graph. The conversation module (120) can then be configured to modify the relevance scores of the accounts based on their corresponding credibility rating. For example, the conversation module (120) can be configured to apply a proportional increase/decrease to the relevance score based on a standard deviation of the accounts credibility score. In another example, the conversation graph (120) can be configured to increase the relevance scores of authoring accounts having a credibility rating exceeding a predefined threshold.

In one or more embodiments of the invention, the conversation module (120) includes functionality to apply a relevance filter to increase the relevance scores of one or more authoring accounts of the conversation graph which are identified in a whitelist of verified accounts. For example, the whitelist of verified accounts can be a list of accounts which are high-profile accounts which are susceptible to impersonation. In this example, celebrity and business accounts would be verified by the messaging platform (100) in order to notify users of the messaging platform (100) that the accounts are authentic. In one or more embodiments of the invention, the conversation module (120) is configured to increase the relevance scores of verified authoring accounts by a predefined amount/percentage.

In one or more embodiments of the invention, the conversation module (120) includes functionality to apply a relevance filter to decrease the relevance scores of one or more authoring accounts of the conversation graph which are identified in a blacklist. Examples of blacklisted accounts can include, but are not limited to, potential spammers, potentially fake accounts, accounts with a potential or history of posting adult content, accounts with a potential or history of posting illegal content, accounts flagged by other users, and/or meeting any other criteria for flagging accounts. For example, accounts with suspicious activity or accounts with a history of violating guidelines of the messaging platform (100), though not currently banned, can be included in the blacklist of accounts. In one or more embodiments of the invention, the conversation module (120) is configured to decrease the relevance scores of blacklisted accounts by a predefined amount/percentage.

In one or more embodiments of the invention, the conversation module (120) includes functionality to apply a relevance filter to entirely exclude messages authored from blacklisted accounts from a conversation graph. Thus, the conversation graph can be configured to designate the messages as non-relevant, irrespective of relevance scoring.

In one or more embodiments of the invention, the conversation module (120) includes functionality to apply a relevance filter to modify the relevance scores of one or more authoring accounts based on context data. For example, the conversation module (120) can be configured to identify a requesting account associated with the request for the relevant conversation. In this example, the requesting account is an account of a user requesting to view the relevant conversation. The user has logged into their account using a web browser of a laptop computer. The user then selected a "view conversation" link within a representation of a single message of a stream of messages displayed in the web browser. Continuing the example, the browser requests the relevant conversation and provides an identifier of the user's account as context data. In one or more embodiments, the conversation module (120) is configured to identify one or more authoring accounts of the conversation graph having a predefined graph relationship with the requesting account. Examples of the predefined graph relationship can include, but are not limited to, friendship, subscription, follower/followee (i.e., accounts followed by requesting account), accounts "favorited" by the requesting account, and/or any other unidirectional or bidirectional relationship between the requesting account and at least one other authoring account of the conversation graph. The conversation module (120) can then, for example, be configured to increase the relevance scores of the accounts having the predefined graph relationship with the requesting account.

In one or more embodiments of the invention, the conversation module (120) includes functionality to apply a relevance filter to modify the relevance scores of one or more authoring accounts based on engagement data. For example, the conversation module (120) can be configured to identify messages having a higher than expected amount of one or more types of engagement. Examples of a type of engagement can include, but are not limited to, replies, favorites, rebroadcasts, shares, views, and/or any other type of measurable interaction with one or more messages. The conversation module (120) can be configured to identify a threshold engagement amount for a given type of engagement, and to compare a measured engagement value with the threshold amount. The measured engagement value can be for one or more messages of the conversation graph. In response to detecting that the measured engagement value of a message exceeds the threshold engagement amount, the conversation module (120) can be configured to increase a relevance score of an authoring account of the message. In this way, the conversation module (120) can be configured to modify the relevance scores of one or more authoring accounts of the conversation graph based on engagement data of one or more messages represented by the conversation graph.

In one or more embodiments of the invention, the conversation module (120) includes functionality to identify relevant participants. Relevant participants may be identified based on a variety of different data. For example, relevant participants may be tagged by the user during composition of the message and may be stored as metadata associated with the message (i.e., not within the message content). In another example, relevant participants may be explicitly tagged within the content of the message, as in the example

discussed above with regard to FIG. 3 (e.g., “@johnqpublic is correct!”). Lastly, in yet another example, relevant participants are inferred based on text or data within the content of the message (e.g., facial recognition, display name matching, etc.). Thus, in one or more embodiments of the invention, the conversation module (120) can be configured to identify a relevant participant based on a character string of text which is not explicitly tagged as an account identifier (e.g., not preceded by the “@” symbol). For example, a user may be identified by the user’s real name. If the user’s real name can be found in display name column 404 of FIG. 4 (described in further detail above), then that user’s account ID can be inserted into relevant participant(s) Account ID 208 field of an appropriate record. The conversation module (120) can be configured to store the relevant participants in a field of a message table (e.g., column 308 of message table 300, discussed above).

In one or more embodiments of the invention, the conversation module (120) includes functionality to apply a relevance filter to modify the relevance scores of one or more authoring accounts based on relevant participants. For example, the conversation module (120) can be configured to increase the relevance scores of authoring accounts which are identified as relevant participants of in a message of a blessed, highly credible, and/or high relevance account.

In one or more embodiments of the invention, the conversation module (120) includes functionality to assign a relevance score to one or more messages of a conversation graph. Thus, each message of the conversation graph is assigned a relevance score (i.e., a message relevance score). The relevance score can be an arbitrary initial score or any other baseline score assigned prior to applying one or more relevance filters. For example, all message relevance scores can be initialized to a value of zero.

In one or more embodiments of the invention, the conversation module (120) includes functionality to apply relevance filters to the conversation graph in order to modify one or more message relevance scores. Relevance filters which are used to modify relevance scores of messages can be referred to as message relevance filters.

In one or more embodiments of the invention, the conversation module (120) includes functionality to apply a relevance filter to increase the message relevance score of a context message and/or a root message of the conversation graph. Similar to the account relevance filters (discussed above), message relevance scores can be modified based on engagement data of the corresponding message, a blessed authoring account, credibility rating of the authoring account, whitelisted authoring accounts, blacklisted authoring accounts, context data, and/or any other criteria for determining relevance and/or importance of a message in the conversation graph.

In one or more embodiments of the invention, the conversation module (120) includes functionality to apply the relevance filters according to a predefined sequence. The sequence of applying relevance filters can be significant in the case of dependency between one or more relevance filters and/or if data required to modify one or more relevance scores otherwise changes over time. For example, if the conversation module (120) is configured to identify blessed accounts based on their relevance scores exceeding a predefined threshold, it may be advantageous to execute one or more other relevance filters (which may modify those relevance scores) prior to designating the blessed accounts. Thus, in one or more embodiments, a credibility relevance filter (i.e., a relevance filter which can modify relevance scores based on a credibility rating), a verified relevance

filter (i.e., a relevance filter which can modify relevance scores based on a whitelist of verified accounts), a relationship graph relevance filter (i.e., a relevance filter which can modify relevance scores based on accounts having a predefined graph relationship with a requesting account), and/or any number of other relevance filters can be executed by the conversation module (120) prior to executing a blessing relevance filter (i.e., a relevance filter which can modify relevance scores based on designating blessed accounts, discussed above). More generally, in one or more embodiments, the conversation module (120) can be configured to identify a volatility ranking of each relevance filter, which indicates a degree of dependency upon the execution of other filters and/or a preferred order of execution. The conversation module (120) can then execute the relevance filters in order of their corresponding volatility ranking.

In another example, engagement data which may be required to apply a relevance filter (as discussed above) accumulates over time. Thus, a relevance filter requiring engagement data can be applied periodically, on-demand, and/or after a fixed time period after the root message is posted. Other data required by one or more relevance filters, such as credibility of an authoring account and verified status of an authoring account, are less likely to change over time. As a result, relevance filters requiring this data can be applied earlier in the sequence and/or aperiodically.

In one or more embodiments of the invention, the conversation module (120) includes functionality to apply the relevance scores by recursively and/or iteratively traversing the conversation graph until convergence. Thus, for example, the conversation module (120) can designate accounts as blessed by traversing the conversation graph and identifying messages which are replied to by blessed accounts. The replied to accounts can then be designated as being blessed, and the process can be repeated until no new replied-to accounts are identified (i.e., convergence).

In one or more embodiments of the invention, the conversation module (120) includes functionality to modify one or more message relevance scores based on one or more account relevance scores. Furthermore, in one or more embodiments of the invention, any one of the account relevance filters can be used by the conversation module (120) to modify one or more message relevance scores by modifying the message relevance scores of all messages posted by an authoring account rather than an account relevance score of the authoring account. In one or more embodiments of the invention, the conversation module (120) is configured to assign a weight to the message relevance score and account relevance score of each message in the conversation graph. For example, the account score can be weighted at 60% while the message relevance score is weighted at 40%. In this way, the conversation module (120) can calculate an aggregate relevance score for each message of the conversation graph.

In one or more embodiments of the invention, the conversation module (120) includes functionality to select a subset of messages of the conversation graph for inclusion in a relevant conversation. The conversation graph (120) can select the messages based on relevance scores of the authoring accounts and/or based on message relevance scores. In one or more embodiments of the invention, the conversation module (120) includes functionality to use a threshold for selecting relevant messages for inclusion in the relevant conversation. The threshold can be a fixed minimum relevance score, a minimum distance from the standard deviation of the relevance scores, a fixed number of the highest

13

relevance scores, and/or any other criteria for modifying the number of messages selected for inclusion in the conversation graph.

In one or more embodiments of the invention, each relevance filter can be either query-dependent or query-independent. Query-independent relevance filters do not require context data. Thus, for example, a query-independent relevance filter can be applied to a conversation graph even prior to receiving a client request for a relevant conversation. In this example, the query-independent filter can be used to prune one or more messages from the conversation graph or to pre-calculate/modify one or more relevance scores. In this way, a pruned conversation graph and/or a pre-calculated relevant conversation can be stored in the conversation repository (148). Similarly, the conversation module (120) can store one or more pre-calculated relevance scores in the conversation repository (148). Subsequently, after receiving context data associated with a request, the conversation module (120) can apply one or more query-dependent relevance scores to the conversation graph in order to identify a relevant conversation on-demand.

In one or more embodiments of the invention, the conversation module (120) is configured to provide the relevant conversation as a list of relevant messages selected from the conversation graph and ordered chronologically by posting time. Alternatively, in one or more embodiments, a reply structure of the messages is maintained and provided to the client (105) (e.g., so the client can display a multi-threaded representation of the relevant conversation).

In one or more embodiments of the invention, the conversation module (120) includes functionality to provide at least a portion of the relevant conversation to a client (e.g., client 105). The portion of the relevant conversation can be selected based on context data, a default requirement, identification of one or more of the message by the client (105), a pagination parameter provided by the client (105), and/or any other criteria for selecting messages of the relevant conversation. Thus, for example, a client may request 10 messages of a relevant conversation, and the conversation module (120) may provide a first 10 messages of the relevant conversation in the absence of any context data. In another example, the client (105) provides a context message in conjunction with the request, and the conversation module (120) provides a window of messages surrounding the context message from the relevant conversation. The window of messages can be a list of messages in chronological order of their posting time to the messaging platform (100) or can include one or more elements of their placement within the reply structure of the conversation graph. Thus, the client (105) can be configured to display a threaded view of the window of messages (e.g., to delineate the separate threads of messages within the window of messages) or to display a non-threaded new (e.g., a list), in accordance with various embodiments of the invention. An example of a threaded view can be some variation of the conversation graph (299) of FIG. 2 displayed in a user interface (UI) of the client (105), with only the messages provided in the window of messages shown. In this example, other messages of the relevant conversation (i.e., those not included in the window of messages), can be represented by a selectable UI element such as an ellipsis. Upon receiving a selection of the selectable UI element, the client (105) can be configured to expand the relevant portion of the relevant conversation by sending a request for a modified window of the relevant conversation (e.g., including a pagination parameter) to the frontend module (110).

14

In one or more embodiments of the invention, the conversation module (120) is configured to receive a pagination parameter and to provide adjacent sections of N (i.e., a predefined number of) messages from the relevant conversation in response to requests including a pagination parameter.

The semantic analysis module (128), in one or more embodiments, is included in the conversation module (120). In other embodiments, the semantic analysis module (128) is a module separate from the conversation module (120). The semantic analysis module (128) includes functionality to identify one or more topics associated with messages, accounts, and/or conversations as described in further detail below.

In one or more embodiments of the invention, the semantic analysis module (128) includes functionality to identify one or more topics associated with messages, accounts, and/or conversations. For example, the semantic analysis module (128) can analyze a message based on frequency of a word, density of a word, connotations of a word, semantic equivalents of a word, sentiments associated with a word, and/or any other criteria for categorizing an account, message, or conversation. In one or more embodiments of the invention, the semantic analysis module (128) includes functionality to modify a relevance score based on a relevance of a message, account, or conversation to a topic. The semantic analysis module (128) can then provide one or more modified relevance scores to the conversation module (120) in response to a request. In one or more embodiments of the invention, the semantic analysis module (128) is configured to identify one or more authors of messages in the conversation with high relevance scores to a topic. In other words, if a conversation has a high relevance score to the topic "basketball" the semantic analysis module (128) may identify one or more professional basketball players with messages in the conversation graph (by virtue of the fact that a basketball player's account may have a high relevance score to the topic "basketball"). In this example, the semantic analysis module (128) then increases the relevance scores of the authoring accounts in proportion with their relevance to the topic.

In one or more embodiments of the invention, as part of applying a relevance filter, the conversation module (120) includes functionality to provide one or more selected topics to the semantic analysis module (128). The semantic analysis module (128) can then modify one or more relevance scores of a conversation graph based on relevance of the messages/accounts/conversation graph to the provided topics. In one or more embodiments, the topics are context data associated with a request for a relevant conversation. The topics can be received from the client (105) in conjunction with a request for a relevant conversation, or can be inferred/selected by the conversation module (120), in accordance with various embodiments of the invention.

In one or more embodiments of the invention, the client (105) is a computing device used to view one or more messages of the messaging platform (100). The client (105) may be configured to download one or more relevant conversations or message streams including any number of broadcast messages from the frontend module (110). The client (105) can then render the broadcast messages for display to a user. The client (105) can be configured to render the messages for display using a web browser, a mobile application, or any other client application, in accordance with various embodiments of the invention.

In one or more embodiments of the invention, the client (105) can be any computing device capable of rendering one

15

or more broadcast messages for display to a user. Examples of a client (105) can include, but are not limited to, a laptop computer, a desktop computer, a server computer, a netbook, a tablet computer, and any other computing device having communication capability with the messaging platform (100). The client (105) can utilize any of a number of advantageous interfaces, including a web-based client, a Short Messaging Service (SMS) interface, an instant messaging interface, an email-based interface, an API function-based interface, and etc. to communicate (via a computer network) with one or more components of the messaging platform (100).

In one or more embodiments of the invention, one or more of the data repositories (message repository 140, connection graph repository 142, stream repository 144, account repository 146, and conversation repository 148) is a database and/or storage service residing on one or more servers. For example, one or more of the data repositories may be implemented as a storage service using service-oriented architecture (SOA) and configured to receive requests for data and to provide requested data to other components of the messaging platform (100). In another example, message repository 140 may include one or more tables in a distributed database management system (DBMS), a clustered database, a standalone flat file, and/or any storage software residing on one or more physical storage devices. Examples of a storage device may include, but are not limited to, a hard disk drive, a solid state drive, and/or other memory device. Any type of database or storage application can be used, in accordance with various embodiments of the invention.

In one or more embodiments of the invention, one or more of the data repositories (message repository 140, connection graph repository 142, stream repository 144, account repository 146, and conversation repository 148) is a separate application or series of applications residing on one or more servers external (and communicatively coupled) to the messaging platform (100). Alternatively, in one or more embodiments of the invention, one or more of the data repositories may be an integrated component of the messaging platform (100) and/or may reside, either partially or entirely, on one or more common hardware devices (e.g., a server).

In one or more embodiments of the invention, routing module 125 includes functionality to receive one or more messages and to store the messages in message repository 140. Routing module 125 may be configured to assign an identifier to the message and to notify graph fanout module 130 of a sender of the message.

In one or more embodiments of the invention, graph fanout module 130 includes functionality to retrieve graph data from connection graph repository 142 and to use the graph data to determine which accounts in messaging platform 100 should receive the message. The graph data, for example, can reflect which accounts in the messaging platform are “following” a particular account and are, therefore, subscribed to receive status messages from the particular account.

In one or more embodiments of the invention, the delivery module 135 includes functionality to receive a list of accounts from graph fanout module 130 and the message identifier generated by routing module 125 and to insert the message identifier into stream data associated with each identified account. The delivery module 135 may then store the message list in stream repository 144. The stream data stored in stream repository 144 can make up one or more streams associated with one or more accounts of messaging platform 100. A stream may be a dynamic list of messages associated with one or more accounts or can reflect any

16

arbitrary organization of messages that is advantageous for the user of an account. Message repository 140 can be used to generate message streams comprising lists of messages for presentation to a user based on various criteria. For example, a message stream can contain a list of messages from accounts that a particular account is associated with (e.g., following, friends with, etc.). Message stream data, in one or more embodiments, is stored in stream repository 144.

In one or more embodiments of the invention, frontend module 110 is a software application or a series of related software applications configured to communicate with external entities (e.g., client 105). Frontend module 110 may include an application programming interface (API 112) and/or any number of other components used for communicating with entities outside of messaging platform 100. The API (112) can include any number of specifications for making requests from and/or providing data to messaging platform 100. For example, a function provided by the API can provide broadcasted messages to a consumer client (e.g., client 105).

In one or more embodiments of the invention, frontend module 110 is configured to use one or more of the data repositories (e.g., message repository 140, connection graph repository 142, stream repository 144, and/or account repository 146) to define streams for serving messages (i.e., stream data) to a user of the account on messaging platform 100. A user can use any client (e.g., client 105) to receive the messages. For example, where the user uses a web-based client to access to messaging platform 100, an API of frontend module 110 can be utilized to define one or more streams and/or to serve the stream data to the client for presentation to the user. Similarly, different forms of message delivery can be handled by different modules in frontend module 110. In one or more embodiments of the invention, the user can specify particular receipt preferences which are implemented by frontend module 110.

FIG. 5 shows a flowchart of a method for using an in-message application. While the various steps in this flowchart are presented and described sequentially, one of ordinary skill will appreciate that some or all of the steps may be executed in different orders and some or all of the steps may be executed in parallel. Further, in one or more embodiments of the invention, one or more of the steps described below may be omitted, repeated, and/or performed in a different order. Accordingly, the specific arrangement of steps shown in FIG. 5 should not be construed as limiting the scope of the invention.

In STEP 500, a request for a relevant conversation is received. In one or more embodiments of the invention, the request is for a predefined number of messages of the relevant conversation. The request can include or be sent in conjunction with any number of context data items. Any number of requests can be received in parallel, and from a variety of different clients (e.g., mobile devices, laptops, netbooks, etc.). For example, the request can be an API call made by the client and including context data parameters.

In STEP 505, a segment of the relevant conversation is identified based on the request. The segment can be identified using the method described with regard to FIG. 7, in accordance with various embodiments of the invention. In one or more embodiments of the invention, a pre-calculated relevant conversation is identified in a repository and retrieved in response to the request.

In STEP 510, the segment of the relevant conversation is provided to the client in response to the request. The segment can include any number of messages of the relevant

17

conversation ordered according to any predefined manner. For example, 20 messages can be included in a list, ordered chronologically by posting time of the messages.

FIG. 6 shows a flowchart of a method for using an in-message application. While the various steps in this flowchart are presented and described sequentially, one of ordinary skill will appreciate that some or all of the steps may be executed in different orders and some or all of the steps may be executed in parallel. Further, in one or more embodiments of the invention, one or more of the steps described below may be omitted, repeated, and/or performed in a different order. Accordingly, the specific arrangement of steps shown in FIG. 6 should not be construed as limiting the scope of the invention.

In STEP 600, a master conversation graph is identified. The master conversation graph includes all messages in a reply structure of messages descending from a root message. Alternatively, in one or more embodiments, one or more of the messages are excluded prior to proceeding to STEP 605 (e.g., based on a relevance filter).

In STEP 605, messages are pruned from the master conversation graph based on context data provided in conjunction with the request. The messages can be pruned based on any criteria for excluding messages from the application of one or more relevance filters, in accordance with various embodiments of the invention. For example, branches of the conversation graph which do not represent a context message, messages from a blessed authoring account, and/or any other identifiable type of message can be excluded from the conversation graph. Excluding messages from the conversation graph, or “pruning” the conversation graph can include tagging the messages for exclusion, creating a new conversation graph including only non-excluded messages, permanently removing the messages from the conversation graph, and/or any other means of identifying a portion of the conversation graph for further analysis. For example, all messages not descended from the context message, or in a direct lineage between and including the root and context messages can be pruned from the conversation graph.

In STEP 610, the pruned conversation graph is used to identify a relevant conversation. For example, the relevant conversation can be identified using the method described with regard to FIG. 7 (below).

FIG. 7 shows a flowchart of a method for using an in-message application. While the various steps in this flowchart are presented and described sequentially, one of ordinary skill will appreciate that some or all of the steps may be executed in different orders and some or all of the steps may be executed in parallel. Further, in one or more embodiments of the invention, one or more of the steps described below may be omitted, repeated, and/or performed in a different order. Accordingly, the specific arrangement of steps shown in FIG. 7 should not be construed as limiting the scope of the invention.

In STEP 700, a request for a segment of a relevant conversation is received from a client. In one or more embodiments of the invention, the request includes a context message. In STEP 705, a conversation graph including a reply structure of messages posted by a set of authoring accounts is identified based on the context data. The conversation graph can be a master conversation graph, or a conversation graph including a preselected subset of messages of a conversation, in accordance with various embodiments of the invention.

In STEP 710, a relevance score is identified for each of the set of authors. The relevance scores can be initialized to a

18

default value (e.g., zero), or can be obtained after modification by applying a relevance filter or some other analysis.

In STEP 715, a relevance filter is applied to the conversation graph in order to modify the relevance scores. In one or more embodiments of the invention, query-independent relevance filters can be applied prior to receiving a request for a relevant conversation. In other words, STEPS 710-720 can be performed prior to receiving the request in STEP 700. In this case, the identified relevant conversation and/or relevance scores can be stored and retrieved later for on-demand processing. For example, after receiving a request, the conversation graph and relevance scores can be retrieved and query-dependent relevance filters can be applied. In one or more embodiments, multiple relevance filters can be applied concurrently. Thus, STEPS 715 and 720 can occur in parallel. In STEP 720, it is determined whether additional relevance filters are designated to be applied to the conversation graph. If it is determined that additional relevance filters exist, the process proceeds to STEP 715. If it is determined that additional relevance filters do not exist, the process proceeds to STEP 725.

In STEP 725, a subset of the messages are selected for inclusion in a relevant conversation based on the relevance scores. The subset of messages can be any portion (or the entirety) of an identified relevant conversation, and can be provided in any format, in accordance with various embodiments. For example, the subset can be provided to the client as a list containing a predefined number of the messages in chronological order of their respective posting times to the messaging platform. Alternatively, in another example, the reply structure of the messages can also be provided to the client. Thus, the client can display the subset of messages represented as a list, or as a multi-threaded tree structure, in accordance with various embodiments of the invention.

The following section describes various examples of the invention. The examples are included to aid in the understanding of the invention but are not intended to limit the scope of the invention.

FIG. 8A depicts example user interface of a mobile device (899) according to one embodiment in which messages are displayed to a user. The particular account for which the user interface is displaying messages is identified by account field 855. Icons depicted at a bottom of the user interface (i.e., 890, 892, 894, and 896) pertain to functions a user can select. Home stream icon 890 can be selected (e.g., by touching a touchscreen over the icon) in order to display a home stream of messages. Settings icon 892 can be selected in order to display a menu of account settings. Messages icon 894 can be selected in order to display private messages (i.e., direct messages) between the account and one or more other accounts. Search icon 896 can be selected in order to search messages based on various user selected criteria.

In the example of FIG. 6A, a home stream of messages is displayed by default. The home stream displays a live feed of messages posted by accounts followed by the user “@keith”. Messages 815 and 830 are displayed in the home stream based on account “@keith” following the account “@brian” as shown in table 400 of FIG. 4. Messages 805, 810, 820, and 825 are also displayed in the home stream based on account “@keith” additionally following the accounts “@ricardo”, “@savepro”, “@debbie”, and “@saveearth” (not shown in FIG. 4). The author of each message, in one embodiment, is identified by the author’s avatar followed by the author’s account ID, which is case insensitive in the current example. Messages 815 and 830 each contain a “View Conversation” link (849, 850) which, in this example, is a hyperlink which indicates that relevant

19

messages including the corresponding message can be viewed by selecting the hyperlink.

As depicted in FIG. 8A, messages **815** and **830** are from account “@brian”. Message **815** is from account “@brian” and is directed to accounts “@adam” and “@joseph” and includes the text “@adam @joseph How can you say that!?”. Message **830** is from account “@brian” and is directed to account “@carrie” and includes the text “@Carrie I love you anyway.” In this example, other relevant messages which are part of the conversation are not displayed in the home stream because they are posted by accounts which are not followed by account “@keith”. In addition, messages which do not pertain to the conversation are also displayed in the home stream (i.e., messages **805**, **810**, **820**, **825**).

Messages **815** and **830** may appear to not make sense and may not provide a user reading the messages with a complete understanding of the interaction. Further, messages **805**, **810**, **820**, and **825** can further confuse a reader since these messages are interspersed between messages **815** and **830** and are not related to the conversation. Hyperlink **849** in message **815** indicates that there may be more messages in the conversation which are not included in the current stream, or are dispersed throughout the stream. In this example, messages **815** and **830** of FIG. 8A are depicted as messages **204** and **212** in conversation graph **299** of FIG. 2.

Continuing the example, the user selects hyperlink **849** in message **815** in order to view relevant messages. Based on the user's selection of hyperlink **849**, message **815** is identified as the context message. Since message **815** corresponds to message **204** of FIG. 2, message **204** is marked as relevant based on its selection as the context message. An application programming interface (API) call is made to the messaging platform requesting 5 messages of a relevant conversation. The API call includes an identifier of the context message.

A conversation module of the messaging platform then obtains a master conversation graph (depicted by FIG. 2) and begins by pruning all branches of the graph not including the context message. Thus, in this example, message **206** of FIG. 2 is removed from the graph for purposes of the analysis. The conversation module then creates a relevance score for each authoring account of the conversation graph. The relevance scores are initialized to zero, as shown in the following table:

Authoring Account	Relevance Score
@adam	0
@joseph	0
@david	0
@carrie	0
@john	0
@ashish	0

Continuing the example, the conversation module applies a series of relevance filters to the conversation graph in a predefined sequence, in order to modify the relevance scores. In this example, the authoring account of the context message is automatically included in the relevant conversation, and is therefore not scored.

In applying the first relevance filter, the conversation module identifies a credibility rating of each authoring account of the conversation graph. All accounts having a

20

credibility rating above 1000 points get a +1 increase in their relevance scores. The following depicts the relevance scores at this point in the example:

Authoring Account	Relevance Score
@adam	1
@joseph	0
@david	1
@carrie	1
@john	0
@ashish	0

The conversation module then applies a second relevance filter. Applying the second relevance filter involves identifying a whitelist of verified users. The conversation graph compares each of the authoring accounts of the conversation graph to the whitelist, and if the authoring account is determined to be verified (i.e., found in the whitelist), then the conversation graph increases the relevance score of that account by +5 points. In this example, “@joseph” and “@carrie” are identified as a verified account. The following depicts the relevance scores at this point in the example:

Authoring Account	Relevance Score
@adam	1
@joseph	5
@david	1
@carrie	6
@john	0
@ashish	0

Continuing the example, the conversation module then applies a third relevance filter. Applying the third relevance filter involves identifying accounts followed by the requesting account (i.e., “@keith”). As represented by the account data table (400) of FIG. 4, the account “@keith” follows only the account “@brian”, which is not represented by a relevance score. Accordingly, the relevance scores are not modified.

Continuing the example, the conversation module then applies a fourth relevance filter. Applying the fourth relevance filter involves designating the authoring account of the context message (i.e., “@brian”) and the authoring account of the root message (i.e., “@adam”) as blessed. Then, the conversation module identifies all messages replied to by blessed authoring accounts and, in turn, designates the authoring accounts of the replied to messages as blessed. This process continues until convergence. In this example, all blessed authoring accounts receive an increase in their corresponding relevance score of +2. The following depicts the relevance scores at this point in the example:

Authoring Account	Relevance Score
@adam	3
@joseph	7
@david	1
@carrie	8
@john	0
@ashish	0

At this point, the conversation module selects the smallest subset of the highest ranking authors which have, in aggre-

21

gate, at least 3 messages in the conversation graph. Only 3 messages are required due to the author of the context message, “@brian” being included in the result set by default. Since “@brian” has two messages in the conversation graph, an additional 3 messages are required.

In this example, the conversation module first includes “@carrie” in the relevant conversation due to her highest relevance score of 8. Carrie has only 1 message in the conversation graph (i.e., less than 4), and so the conversation module proceeds to select the next highest relevance score. In this example, “@joseph” has the next highest relevance score and is therefore selected for inclusion in the relevant conversation. After including “@joseph”, the relevant conversation includes 2 messages. The conversation module determines that at least 1 more message is required and, thus, selects the next highest relevance scoring account, “@adam”, for inclusion in the relevant conversation. At this point the relevant conversation includes messages 835, 840, 815, 845, and 830 of FIG. 8B which correspond to messages 200, 202, 204, 208, and 212 shown in FIG. 2.

FIG. 8B depicts an example user interface of the mobile device (899) which displays the relevant conversation including messages 835, 840, 815, 845, and 830 in response to selection of hyperlink 849 in message 815 of FIG. 8A. The messages displayed in FIG. 8A are based on accounts “Keith” (i.e., account “@keith”) follows, while the messages displayed in FIG. 8B are messages included in the relevant conversation generated based on an embodiment of the method depicted in FIG. 7 (discussed above). As depicted in FIG. 8B, messages 815 and 830 are augmented by the display of messages 835, 840, and 845. Since messages determined to be part of the relevant conversation are depicted in FIG. 8B, the user associated with the account “@keith” can gain a better understanding of the exchange.

Embodiments of the invention may be implemented on virtually any type of computer regardless of the platform being used. For example, as shown in FIG. 9, a computer system (900) includes one or more processor(s) (902) (such as a central processing unit (CPU), integrated circuit, hardware processor, etc.), associated memory (904) (e.g., RAM, cache memory, flash memory, etc.), a storage device (906) (e.g., a hard disk, an optical drive such as a compact disk drive or digital video disk (DVD) drive, a flash memory stick, etc.), a network adapter (918), and numerous other elements and functionalities typical of today’s computers (not shown). One or more components of computer system 900 may be communicatively connected by a bus (916). Computer system 900 may also include input means, such as a keyboard (908), a mouse (910), or a microphone (not shown). Further, computer system 900 may include output means, such as a monitor (912) (e.g., a liquid crystal display (LCD), a plasma display, or cathode ray tube (CRT) monitor). Computer system 900 may be connected to a network (914) (e.g., a local area network (LAN), a wide area network (WAN) such as the Internet, or any other type of network) via network adapter 918. Those skilled in the art will appreciate that many different types of computer systems exist, and the aforementioned input and output means may take other forms. Generally speaking, computer system 900 includes at least the minimal processing, input, and/or output means necessary to practice embodiments of the invention.

Further, in one or more embodiments of the invention, one or more elements of the aforementioned computer system (900) may be located at a remote location and connected to the other elements over a network. Further, embodiments of the invention may be implemented on a distributed system having a plurality of nodes, where each portion of the

22

invention (e.g., conversation module 120, semantic analysis module 122, frontend module 110, routing module 125, etc. of FIG. 1, discussed above) may be located on a different node within the distributed system. In one embodiment of the invention, the node corresponds to a computer system. Alternatively, the node may correspond to a processor with associated physical memory. The node may alternatively correspond to a processor or micro-core of a processor with shared memory and/or resources. Further, software instructions in the form of computer readable program code to perform embodiments of the invention may be stored, temporarily or permanently, on a non-transitory computer readable storage medium, such as a compact disc (CD), a diskette, a flash memory device, a solid state drive, a primary storage device, or any other tangible computer readable storage device.

One or more embodiments of the invention have one or more of the following advantages. By identifying relevant messages in a conversation graph and providing the relevant messages to a client, it may be possible to provide a concise and relevant view of a conversation to a user. Thus, in one or more embodiments, the user can be spared from viewing less relevant messages, and can view a relevant subset of a relevant conversation including messages from a variety of different accounts based on context data.

While the invention has been described with respect to a limited number of embodiments, those skilled in the art, having benefit of this disclosure, will appreciate that other embodiments can be devised which do not depart from the scope of the invention as disclosed herein. Accordingly, the scope of the invention should be limited only by the attached claims.

What is claimed is:

1. A method for identifying a relevant conversation, comprising:
 - identifying a conversation graph comprising identifiers of a plurality of messages posted by a plurality of authoring accounts;
 - identifying a plurality of relevance scores, wherein each of the plurality of relevance scores corresponds to an authoring account of the plurality of authoring accounts;
 - applying, by a computer processor, a first relevance filter to the conversation graph in order to modify the plurality of relevance scores;
 - applying, by the computer processor and after applying the first relevance filter, a second relevance filter to the conversation graph in order to further modify the plurality of relevance scores;
 - selecting, based on the plurality of relevance scores and after applying the second relevance filter, a subset of messages of the conversation graph for inclusion in the relevant conversation; and
 - providing at least a portion of the relevant conversation to a client device.
2. The method of claim 1 further comprising:
 - identifying a context message of the conversation graph;
 - applying a third relevance filter to the conversation graph by:
 - identifying an authoring account of the context message;
 - identifying a predefined graph relationship between the authoring account of the context message and at least one of the plurality of authoring accounts; and
 - further modifying the plurality of relevance scores based on the predefined graph relationship.

23

3. The method of claim 2, wherein the context message is one selected from a group consisting of a message selected by a user, a message comprising a predefined keyword, a message comprising a predefined topic, a message returned in response to a search request, and a message associated with an advertisement.

4. The method of claim 1, further comprising:

storing the relevant conversation in a conversation repository;

receiving a request for a first segment of the relevant conversation from the client device, wherein the request for the first segment comprises an identifier of a first context message of the conversation graph; and

identifying a portion of the relevant conversation based on the first context message, wherein providing the portion of the relevant conversation to the client device is in response to the request for the first segment, and wherein the first segment comprises a predefined window of message identifiers surrounding the identifier of the first context message within the relevant conversation.

5. The method of claim 1, wherein applying the first relevance filter further comprises:

identifying credibility scores of the plurality of authoring accounts; and

modifying the plurality of relevance scores based on the credibility scores.

6. The method of claim 5, further comprising:

identifying a history of posted messages for each of the plurality of authoring accounts; and

calculating the credibility scores for each of the plurality of authoring accounts based on the history of posted messages.

7. The method of claim 5, wherein applying the second relevance filter comprises:

determining that a subset of the plurality of authoring accounts are verified accounts; and

increasing the relevance scores of the verified accounts.

8. The method of claim 7, further comprising:

receiving, from a client computing device, a request for the relevant conversation, wherein providing the relevant conversation is in response to the request; and applying a third relevance filter by:

identifying a requesting account of the request, and

identifying a predefined graph relationship between the requesting account and an authoring account of the plurality of authoring accounts; and

increasing the relevance score of the authoring account based on the predefined graph relationship.

9. The method of claim 8, wherein the predefined graph relationship is one selected from a group consisting of a friendship, a follower relationship, a subscription, and a fan relationship.

10. The method of claim 8, further comprising:

calculating an engagement metric for each message of the conversation graph, wherein the engagement metric represents a type of engagement with the message by other accounts of a messaging platform;

identifying a message relevance score for each of the plurality of messages; and

applying a fourth relevance filter by modifying the message relevance score of each message based on the engagement metric of the message, wherein selecting the subset of messages of the conversation graph for inclusion in the relevant conversation is further based on the message relevance score of each message.

24

11. The method of claim 1, wherein applying the first relevance filter comprises:

identifying a root message of the conversation graph; designating an authoring account of the root message as being blessed;

iteratively designating, among the plurality of authoring accounts, authoring accounts of messages replied to by blessed authoring accounts as being blessed; and increasing the relevance scores of messages authored by blessed authoring accounts.

12. The method of claim 1, wherein applying the first relevance filter comprises recursively traversing messages of the conversation graph.

13. The method of claim 1, wherein the relevant conversation comprises identifiers of the subset of messages ordered chronologically by posting time.

14. The method of claim 1, wherein identifying the conversation graph comprises:

obtaining an identifier of a context message within a master conversation graph; and

pruning the master conversation graph to generate the conversation graph, wherein pruning the master conversation graph comprises pruning messages not residing in a branch of the context message.

15. A system for identifying a relevant conversation, comprising:

a computer processor; and

a conversation module executing on the computer processor and configured to:

identify a conversation graph comprising identifiers of a plurality of messages posted by a plurality of authoring accounts;

identify a plurality of relevance scores, wherein each of the plurality of relevance scores corresponds to an authoring account of the plurality of authoring accounts;

apply a first relevance filter to the conversation graph in order to modify the plurality of relevance scores;

apply, after applying the first relevance filter, a second relevance filter to the conversation graph in order to further modify the plurality of relevance scores;

select, based on the plurality of relevance scores and after applying the second relevance filter, a subset of messages of the conversation graph for inclusion in the relevant conversation; and

provide at least a portion of the relevant conversation to a client device.

16. The system of claim 15, the conversation module further configured to:

identify a context message of the conversation graph; apply a third relevance filter to the conversation graph by: identifying an authoring account of the context message;

identifying a predefined graph relationship between the authoring account of the context message and at least one of the plurality of authoring accounts; and further modifying the plurality of relevance scores based on the predefined graph relationship.

17. The system of claim 15, the conversation module further configured to:

receive, from a client computing device, a request for the relevant conversation, wherein providing the relevant conversation is in response to the request; and

wherein applying the first relevance filter comprises: identifying a requesting account of the request, and identifying a predefined graph relationship between

25

the requesting account and an authoring account of the plurality of authoring accounts; and increasing the relevance score of the authoring account based on the predefined graph relationship.

18. The system of claim 15, wherein applying the first relevance filter comprises:

identifying a root message of the conversation graph; designating an authoring account of the root message as being blessed;

recursively designating, among the plurality of authoring accounts, authoring accounts of messages replied to by blessed authoring accounts as being blessed; and increasing the relevance score of messages authored by blessed authoring accounts.

19. The system of claim 15, wherein the relevant conversation is a chronologically ordered list, and wherein the system further comprises:

a conversation repository configured to store the relevant conversation; and

wherein the conversation module is further configured to: store the relevant conversation in the conversation repository;

receive a request for a first segment of the relevant conversation from the client device, wherein the request for the first segment comprises an identifier of a first context message of the conversation graph; and

identify a portion of the relevant conversation based on the first context message, wherein providing the portion of the relevant conversation to the client device is in response to the request for the first segment, and wherein the first segment comprises a predefined window of message identifiers surrounding the identifier of the first context message within the relevant conversation.

20. The system of claim 19, wherein the conversation module is further configured to:

receive, from a different client device, a request for a second segment of the relevant conversation, wherein the request for the second segment comprises an identifier of a second context message;

identify the second segment of the relevant conversation based on the second context message, wherein the second segment comprises a predefined window of message identifiers surrounding the identifier of the second context message within the relevant conversation; and

provide the second segment to the different client device in response to the request for the second segment.

21. A non-transitory computer readable medium comprising a plurality of instructions for identifying a relevant conversation, the plurality of instructions, when executed by a computer processor, comprising functionality to:

identify a conversation graph comprising identifiers of a plurality of messages posted by a plurality of authoring accounts;

identify a plurality of relevance scores, wherein each of the plurality of relevance scores corresponds to an authoring account of the plurality of authoring accounts;

apply a first relevance filter to the conversation graph in order to modify the plurality of relevance scores;

26

apply, after applying the first relevance filter, a second relevance filter to the conversation graph in order to further modify the plurality of relevance scores;

select, based on the plurality of relevance scores and after applying the second relevance filter, a subset of messages of the conversation graph for inclusion in the relevant conversation; and

provide at least a portion of the relevant conversation to a client device.

22. The non-transitory computer readable medium of claim 21 wherein the instructions further comprise:

identifying a context message of the conversation graph; applying a third relevance filter to the conversation graph by: identifying an authoring account of the context message;

identifying a predefined graph relationship between the authoring account of the context message and at least one of the plurality of authoring accounts; and further modifying the plurality of relevance scores based on the predefined graph relationship.

23. The non-transitory computer readable medium of claim 21, wherein the instructions further comprise:

storing the relevant conversation in a conversation repository;

receiving a request for a first segment of the relevant conversation from the client device, wherein the request for the first segment comprises an identifier of a first context message of the conversation graph; and identifying a portion of the relevant conversation based on the first context message, wherein providing the portion of the relevant conversation to the client device is in response to the request for the first segment, and wherein the first segment comprises a predefined window of message identifiers surrounding the identifier of the first context message within the relevant conversation.

24. The non-transitory computer readable medium of claim 23, wherein the instructions further comprise:

identifying a history of posted messages for each of the plurality of authoring accounts; and

calculating the credibility scores for each of the plurality of authoring accounts based on the history of posted messages.

25. The non-transitory computer readable medium of claim 21, wherein applying the first relevance filter comprises:

identifying a root message of the conversation graph; designating an authoring account of the root message as being blessed;

iteratively designating, among the plurality of authoring accounts, authoring accounts of messages replied to by blessed authoring accounts as being blessed; and increasing the relevance scores of messages authored by blessed authoring accounts.

26. The non-transitory computer readable medium of claim 21, wherein identifying the conversation graph comprises:

obtaining an identifier of a context message within a master conversation graph; and

pruning the master conversation graph to generate the conversation graph, wherein pruning the master conversation graph comprises pruning messages not residing in a branch of the context message.

* * * * *